601.436/636 Algorithmic Game Theory **Lecturer:** Michael Dinitz
Topic: Facility Location Game **Date:** 3/3/22

12.1 Introduction

Today we're going to discuss a more complex game (the facility location game) which we will analyze via smoothness. Let's first recall what was happening with smooth games.

12.2 Smooth Games

Let's remember our definition of smooth games:

Definition 12.2.1. *A cost-minimization game with objective function* $cost : S \to \mathbb{R}$ *is* (λ, μ)*-smooth if* $cost(s) \leq \sum_{i=1}^{k} C_i(s)$ *for all* $s \in S$ *and*

$$\sum_{i=1}^{k} C_i(s_{-i}, s'_i) \leq \lambda \cdot cost(s') + \mu \cdot cost(s)$$

for all $s, s' \in S$.

Last class we proved that every (λ, μ)-smooth game has price of total anarchy at most $\lambda/(1-\mu)$, and we proved that atomic routing games with affine cost functions are $(5/3, 1/3)$-smooth.

12.2.1 Utility Maximization

We're going to want to talk about smooth utility maximization games, rather than cost minimization games. Everything works exactly as you would expect, so I'll just state the main definitions and results that we'll use.

Definition 12.2.2. *A utility maximization game with global objective function* $V : S \to \mathbb{R}$ *(usually called the value) is* (λ, μ)*-smooth if*

- $V(s) \geq \sum_{i=1}^{k} u_i(s)$ *for all* $s \in S$, *and*
- $\sum_{i=1}^{k} u_i(s_{-i}, s'_i) \geq \lambda \cdot V(s') - \mu \cdot V(s)$ *for all* $s, s' \in S$.

Since we want to keep thinking about the *price* of anarchy/total anarchy, for a utility maximization game let's define the price of total anarchy as

$$\frac{\max_{s \in S} V(s)}{\min_{\sigma \in CCE} \mathbf{E}_{s \sim \sigma}[V(s)]}.$$

In other words, we switch the numerator and denominator compared to cost-minimization games (and note that the worst CCE is now a min over CCEs rather than a max) so that this price is always at least 1.

Theorem 12.2.3. *The price of total anarchy in any (λ, μ)-smooth utility-maximization game is at most $\frac{1+\mu}{\lambda}$.*

Proof. Let σ be a CCE and let s^* be the optimal profile. Then

$$\begin{aligned}\mathop{\mathbf{E}}_{s\sim\sigma}[V(s)] &\geq \mathop{\mathbf{E}}_{s\sim\sigma}\left[\sum_{i=1}^{k} u_i(s)\right] = \sum_{i=1}^{k} \mathop{\mathbf{E}}_{s\sim\sigma}[u_i(s)] \geq \sum_{i=1}^{k} \mathop{\mathbf{E}}_{s\sim\sigma}[u_i(s_{-i}, s_i^*)] = \mathop{\mathbf{E}}_{s\sim\sigma}\left[\sum_{i=1}^{k} u_i(s_{-i}, s_i^*)\right] \\ &\geq \mathop{\mathbf{E}}_{s\sim\sigma}[\lambda \cdot V(s^*) - \mu \cdot V(s)] = \lambda \cdot V(s^*) - \mu \cdot \mathop{\mathbf{E}}_{s\sim\sigma}[V(s)].\end{aligned}$$

Rearranging yields

$$V(s^*) \leq \frac{1+\mu}{\lambda} \mathop{\mathbf{E}}_{s\sim\sigma}[V(s)]$$

as claimed. □

12.3 Facility Location Game

For the rest of today we're going to talk about another game (not routing) which we will also prove is smooth, but which is a bit more complex and so is good practice to see how to use these ideas. This called the "location game", or the "facility location game", or (if we want to be complete accurate and descriptive) "competitive facility location with price-taking markets and profit-maximizing firms".

12.3.1 Definition

- There is a set F of possible "locations".
- There are k players. Player i has as its strategy set some $F_i \subseteq F$. Think of these as "places where player i might build a facility". We'll slightly abuse notation and let $\emptyset \in F_i$, so all players also have the ability to build "nowhere".
- There is a set M of *markets*. Each $j \in M$ has some value $v_j \geq 0$. Think of this as how much the customer at market j is willing to pay for the service that the facilities provide.
- For each $x \in F$ and $j \in M$, there is some cost $c_{xj} \geq 0$ of serving market j from location x.

12.3.1.1 Utilities: example

The utilities are going to end up looking a bit complicated and weird formally, but are pretty intuitive if we work through an example. They're basically the profits that a player gets by putting their facility at some location and trying to serve the markets using their facility. Let's motivate this by an example. We'll let $F = [3]$ and $M = [2]$.

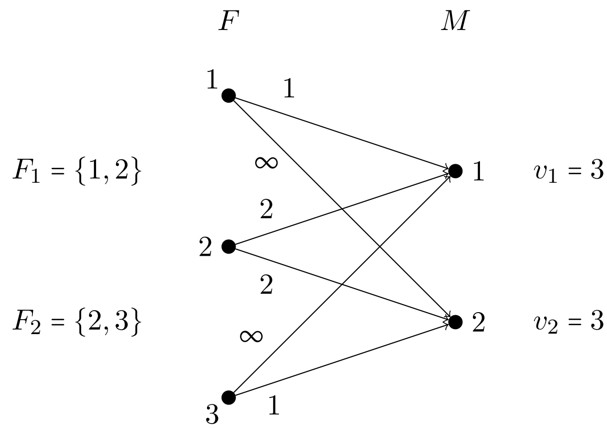

Each player gets to choose a location for their facility, and can also choose a price to charge each market. Each market will choose the player offering it the cheapest price (if it's below their value), and then gives that amount of money to the player, who has to pay the cost for servicing the market from their facility and gets to keep the leftover as profit.

So, for example, suppose that $s_1 = 1$ and $s_2 = 3$ (player 1 chooses location 1 and player 2 chooses location 3). Then player 2 would have to pay an infinite amount to service market 1, so the price it offers market 1 is ∞ (which player 1 will never choose since it's larger than $v_1 = 3$). Player 1, on the other hand, can charge 3 to market 1, and market 1 will take it. This gives profit (utility) of $3 - c_{11} = 3 - 1 = 2$ to player 1. And note that the same thing happens symmetrically with player 2 and market 2, so under this strategy profile each player gets utility 2.

But now suppose that player 1 changes to instead play location 2. Then player 1 can still charge 3 to market 1, but if player 2 tries to charge 3 to market 2 then player 1 can charge less than 3 to undercut them. The lowest price that player 1 can afford to charge market 2 is 2, in order to cover the cost of servicing, and the lowest price that player 2 can afford to charge is 1 (to cover the cost of servicing). So player 2 can charge 2 (or really $2 - \epsilon$) to market 2, and player 1 won't bother. So in this new strategy profile, the utility of player 1 is $3 - 2 = 1$ and the utility of player 2 is $2 - 1 = 1$.

12.3.1.2 Utilities: Definition

Let's formally define utilities now that we've built up some intuition. To simplify notation, let's assume that $c_{xj} \leq v_j$ for all $x \in F$ and $j \in M$. This is without loss of generality, since if we change any service cost larger than v_j to v_j then nothing will change (e.g., in our example changing the ∞ costs to 3 doesn't actually change anything).

Let $s \in S$ be a strategy profile. Let's start with a simple observation: player i at location s_i can get profit from market j only if it's the "closest" (minimum service cost): $c_{s_i j} \leq c_{s_x j}$ for all $x \in [k]$. If it *is* the closest, then the highest price it can charge is the second-smallest cost: $p_{ij}(s) = \min_{x \neq i} c_{s_x j}$. So then player i will get profit from market j equal to

$$\pi_{ij}(s) = \begin{cases} p_{ij}(s) - c_{s_i j} & \text{if } c_{s_i j} \leq c_{s_x j} \text{ for all } x \in [k] \\ 0 & \text{otherwise} \end{cases}$$

Putting all this together, we get that the utility of player i under strategy profile s is

$$u_i(s) = \sum_{j \in M} \pi_{ij}(s).$$

12.3.1.3 Global Value

Now let's define the global value function $V : S \to \mathbb{R}$. We don't want to just use the sum of the utilities, since that looks unfair to the markets (even though they're not technically players in the game). Instead, we'll use the sum of the utilities *including* the markets. Confusingly, this is usually called "social surplus" or "social welfare" even though we're adding in agents which are not actually players.

Given a strategy profile s, let $f(j)$ be the player who serves market j, and let $p_{f(j)j}(s)$ be the price charged to market j. Then the sum of the player utilities is

$$\sum_{i=1}^{k} u_i(s) = \sum_{i=1}^{k} \sum_{j \in M} \pi_{ij}(s) = \sum_{j \in M} \pi_{f(j)j}(s) = \sum_{j \in M} \left(p_{f(j)j}(s) - c_{s_{f(j)}j} \right).$$

The "utility" of a market is the natural thing you'd expect: the value that it puts on getting the service minus the price that it pays. Then the sum of market utilities is

$$\sum_{j \in M} \left(v_j - p_{f(j)j}(s) \right).$$

Adding these together gives us our overall definition of value:

$$V(s) = \sum_{j \in M} \left(p_{f(j)j}(s) - c_{s_{f(j)}j} \right) + \sum_{j \in M} \left(v_j - p_{f(j)j}(s) \right) = \sum_{j \in M} \left(v_j - c_{s_{f(j)}j} \right).$$

Note that the prices don't end up mattering to the social welfare/surplus! Prices redistributed the social welfare between the players and the markets, but the "total happiness" is independent of the actual prices.

12.3.2 Smoothness

It turns out that the location game is actually a potential game, so there is always a pure Nash. We're not going to prove this, but it's a good exercise to do (it's an exercise in Roughgarden's book as well). Instead, we're going to bound the price of total anarchy by proving the following theorem.

Theorem 12.3.1. *The location game is* $(1,1)$*-smooth.*

This implies by Theorem 12.2.3 that the price of total anarchy (and thus also the price of anarchy and pure price of anarchy) is at most 2.

To prove this, we're going to establish three properties. Once we have them, it will be simple to prove smoothness.

Property 1: $\sum_{i=1}^{k} u_i(s) \le V(s)$ for all $s \in S$. This is obviously true since we defined $V(s)$ to be the sum of player utilities plus market utilities, and market utilities are never negative. Note that we need this property for the definition of smoothness.

Property 2: $u_i(s) = V(s) - V(s_{-i})$ for all $i \in [k]$ and $s \in S$. In other words, utility is equal to surplus created. We're abusing notation here slightly, but hopefully it's obvious what we mean by $V(s_{-i})$: consider the same game but where i is not a player, and let s_{-i} be the strategy profile of that game obtained by removing s_i from s, and let $V(s_{-i})$ denote the social surplus of s_{-i} in that game.

For each market j let $g(j)$ be the closest location in s (the closest location to j in which at least one player places a facility) and let $g_{-i}(j)$ be the same thing for s_{-i}. Let's start from the right hand side of what we're trying to prove.

$$V(s) - V(s_{-i}) = \sum_{j \in M} \left((v_j - c_{g(j)j}) - (v_j - c_{g_{-i}(j)j}) \right) = \sum_{j \in M} (c_{g_{-i}(j)j} - c_{g(j)j})$$

What is $c_{g_{-i}(j)j} - c_{g(j)j}$? If $g(j) \neq s_i$ then removing player i makes no difference, and so then $c_{g_{-i}(j)j} - c_{g(j)j} = 0$. On the other hand, if $g(j) = s_i$, then i is the closest player j in s and whatever player is on location $g_{-i}(j)$ is the second closest player. In other words: if in s player i is not the closest player to j then this term is 0, and otherwise it's the second-closest minus the closest. This is precisely the definition we had of profit! So $c_{g_{-i}(j)j} - c_{g(j)j} = \pi_{ij}(s)$. Thus

$$V(s) - V(s_{-i}) = \sum_{j \in M} \pi_{ij}(s) = u_i(s)$$

Property 3: V is *monotone* and *submodular*.

By monotone, we mean that adding more players cannot decrease the social welfare. Slightly more formally, $V(s_{-i}) \le V(s)$ for all $s \in S$ (and note that by induction this extends to any subset of s, not just removing one player). This is easy to see from property 2, or from the fact that adding more players does not affect the v values and can only decrease the c values in the definition of V.

To prove (and define) submodularity, I'm going to slightly abuse notation and for a subset $s \subseteq F$ let $V(s)$ denote the social welfare given by these locations. I can do this because the definition of welfare/surplus is actually only about the locations that players choose, not about the players themselves. I claim that V is *submodular*:

$$V(s' \cup \{x\}) - V(s') \le V(s \cup \{x\}) - V(s)$$

for all $x \in F$ and $s \subseteq s' \subseteq F$. In other words, V has decreasing marginal gains: adding x helps more for s than for any set containing s.

To prove this, first note that by property 2 this is equivalent to proving that $u_x(s' \cup \{x\}) \le u_x(s \cup \{x\})$. Now by definition, this is equivalent to proving $\sum_{j \in M} \pi_{xj}(s' \cup \{x\}) \le \sum_{j \in M} \pi_{xj}(s \cup \{x\})$. Let's prove this inequality for each term.

Consider some $j \in M$. If $\pi_{xj}(s' \cup \{x\}) = 0$ then clearly the inequality is true. Otherwise, if $\pi_{xj}(s' \cup \{x\}) > 0$ the x is the closest to j in $s' \cup \{x\}$, and thus is also the closest to j in $s \cup \{x\}$.

So the profit gained by x is the price it can charge (the cost of servicing from the second-closest) minus its servicing cost (the same in both $s \cup \{x\}$ and $s' \cup \{x\}$). Since $s \subseteq s'$, the second-closest in s' is at least as close to j as the second-closest in s, and thus $p_{xj}(s' \cup \{x\}) \le p_{xj}(s \cup \{x\})$. Hence $\pi_{xj}(s' \cup \{x\}) \le \pi_{xj}(s \cup \{x\})$, which implies that $\sum_{j \in M} \pi_{xj}(s' \cup \{x\}) \le \sum_{j \in M} \pi_{xj}(s \cup \{x\})$ and thus that $u_x(s' \cup \{x\}) \le u_x(s \cup \{x\})$ as required.

Final Proof. Let's use these properties to actually prove smoothness. Let's $s, s' \in S$, and we'll also interpret them as sets of locations ($s, s' \subseteq F$)

$$\begin{aligned}
\sum_{i=1}^{k} u_i(s_{-i}, s'_i) &= \sum_{i=1}^{k} \big(V(s_{-i}, s'_i) - V(s_{-i})\big) && \text{(property 2)}\\
&\ge \sum_{i=1}^{k} \big(V(s \cup \{s'_1, s'_2, \dots, s'_i\}) - V(s \cup \{s'_1, s'_2, \dots, s'_{i-1}\})\big) && \text{(submodularity)}\\
&= V(s \cup s') - V(s) && \text{(telescoping)}\\
&\ge V(s') - V(s) && \text{(monotonicity)}
\end{aligned}$$

Together with property 1, this implies $(1,1)$-smoothness.

12.4 Utility Games

Like we've seen with other games, we can generalize this argument to a wider class of games, known as *monotone utility games* [Vet02]. In a monotone utility game each player has a set of strategies A_i, and we let $A = \cup_{i=1}^{k} A_i$ (the "locations") and $S = A_1 \times \dots A_k$ be the strategy profiles. There is a global value function $V : 2^A \to \mathbb{R}$ (i.e., the value function takes every subset $T \subseteq A$ to a real number $V(T)$). Note that, like the with location game, this notion of value is a function only of the locations chosen, not of prices, who chose each location, or anything else. For a strategy profile $s \in A_1 \times A_2 \times \dots \times A_k$, we can slightly abuse notation and define $V(s) = V(\cup_{i=1}^{k} \{s_i\})$. Such a game is a *monotone utility game* if it satisfies the following properties.

1. $V(s) \ge \sum_{i=1}^{k} u_i(s)$ (recall that this is necessary for smoothness).
2. V is submodular: $V(T \cup \{x\}) - V(T) \ge V(T' \cup \{x\}) - V(T')$ for all $T \subseteq T' \subseteq A$ and $x \in A$.
3. V is monotone: $V(T) \le V(T')$ for all $T \subseteq T' \subseteq A$.
4. The utility of a player is *at least* the value created: $u_i(s) \ge V(s) - V(s_{-i})$ for all $i \in [k]$ and $s \in S$. If this inequality is actually an equality (like in the location game) then this is known as a *basic* monotone utility game.

It's easy to see that the facility location game is a basic monotone utility game, since we proved all of the above properties.

It's also easy to see that with these properties the smoothness argument is the same as before (and we don't even need the basic property).

Theorem 12.4.1. *Every monotone utility game is* $(1,1)$*-smooth.*

Proof. Let $s, s' \in S$ be two strategy profiles.

$$\begin{aligned}
\sum_{i=1}^{k} u_i(s_{-i}, s'_i) &\geq \sum_{i=1}^{k} \left(V(s_{-i}, s'_i) - V(s_{-i}) \right) && \text{(property 4)} \\
&\geq \sum_{i=1}^{k} \left(V(s \cup \{s'_1, s'_2, \ldots, s'_i\}) - V(s \cup \{s'_1, s'_2, \ldots, s'_{i-1}\}) \right) && \text{(submodularity)} \\
&= V(s \cup s') - V(s) && \text{(telescoping)} \\
&\geq V(s') - V(s) && \text{(monotonicity)}
\end{aligned}$$

The other requirement for smoothness is from property 1. □

Note that "smoothness" had not invented when Vetta wrote [Vet02]: like with atomic routing games, the original proofs were about pure Nash equilibria and it was only later shown that the proof actually gave smoothness, also implying bounds on the price of total anarchy.

References

[Vet02] Adrian Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, page 416. IEEE Computer Society, 2002.